Agenda

Department of Energy Centers of Excellence Performance Portability Meeting

April 18–21, 2016,
Glendale, AZ

MONDAY, APRIL 18

5:30p	Registration, Reception (until 7:30)

TUESDAY, APRIL 19, TECHNICAL SESSIONS

7:30a	Registration, Coffee, Pastries
Overviews: Rob Neely (LLNL), Chair	
8:15	Welcome/Kickoff
8:30	Tjerk Straatsma, ORNL: Summit COE/CAAR Overview
8:40	Jack Deslippe, LBL: NERSC-8 COE/NESAP Overview
8:50	Rob Neely, LLNL: Sierra COE Overview
9:00	Hai Ah Nam, LANL: Trinity COE Multi-Lab Overview
9:15	Kalyan Kumaran, ANL: ANL COE Overview
9:25	Nick Romero, ANL: HPCOR Workshop Recap
9:35	Bert Still, LLNL/Multi-Lab: ECP Application Overview and Criteria
9:50	BREAK
NDA Sessions (individuals or institutions must be covered under proper non-disclosure agreements)	
10:05	Intel NDA Session
11:10	BREAK
11:20	NVIDIA NDA Session

TUESDAY SESSIONS, continued	
12:25	LUNCH (on your own)
Applications/Optimizations/Algorithms: Rebecca Hartman-Baker (LBL), Chair	
1:45	Jae-Seung Yeom, LLNL: Data-Dependent Performance Modeling of Linear Solvers for Sparse Matrices
2:00	Charles Ferenbaugh, LANL: Coarse versus Fine-Level Threading in the PENNANT Mini-App
2:15	Scott Parker, ANL: Performance Optimization and Portability of the Nekbone Mini-App
2:30	Kris Garrett, LANL: A First Look at Optimizing Performance on the KNL
2:45	Vitali Morozov, ANL: Portability of HACC—a Highly Tuned Cosmology Application
3:00	BREAK (snacks provided)
3:15	Kristopher Keipert, ANL: Experiences and Challenges while Modernizing GAMESS for Theta and Aurora
3:30	Steve Rennich, NVIDIA: GPU Performance Optimization of the Sweep Operation in Kripke
3:45	Balint Joo, JLab/ANL/LBL: Experiences and Challenges for Performance Portability in Lattice QCD
4:00	Alvaro Vazquez-Mayagoitia, ANL: Many-Core and GPU Developments in the Parallel ELectronic Structure Infrastructure Library (ELSI)
4:10	BREAK
Performance Portable Abstractions: Hai Ah Nam (LANL), Chair	
4:40	Tan Nguyen, LBL: Portable Data Locality Management with High-Level Programming Abstractions
4:55	Jeff Vetter, ORNL: Understanding Portability of a High-Level Programming Model on Diverse HPC Architectures
5:15	Christian Trott, SNL: Kokkos—Performance Portability Today
5:35	Rich Hornung, LLNL: The RAJA Encapsulation Model for Architecture Portability
5:55	Arpith Jacob, IBM: Towards Performance Portable GPU Programming with RAJA

TUESDAY SESSIONS, continued	
6:15	ADJOURN (dinner on your own)

WEDNESDAY, APRIL 20, TECHNICAL SESSIONS

7:30a	Coffee, Pastries
8:15	Opening Remarks, Welcome, Recap of Day 1, Overview of Day 2
Managing the Memory Hierarchy: Nick Romero (ANL), Chair	
8:20	David Poliakoff, LLNL: Copy Hiding Application Interface (CHAI)—Hiding Data Motion for Performance Portability
8:30	Nikolai Sakharnykh, NVIDIA: Harnessing Performance of Geometric Multi-Grid Methods by Using LOC and TOC Architectures
8:45	Fabian Delalondre, ANL: Leveraging Heterogeneous Systems and Deep Memory Hierarchies for Brain Tissue Modeling
9:05	Luiz DeRose, Cray: Cray's Programming Environment for Portable Performance and Programmability on Systems with High-Bandwidth Memory
9:20	Ian Karlin, LLNL/Multi-Lab: Quad Lab Proposal of Fundamental Cross Architecture Multi-Level Memory Support
Application Experience with Performance Portable Abstractions: Tjerk Straatsma (ORNL), Chair	
9:40	Changhoan Kim, IBM: An Abstraction for Unstructured Mesh Problems
9:55	Adam Kunen, LLNL: Nested Loop RAJA for Performance Portability
10:10	Stan Moore, SNL: Obtaining Threading Performance Portability in SPARTA Using Kokkos
10:25	BREAK

WEDNESDAY SESSIONS, continued	
10:55	David Beckingsale, LLNL: Lightweight Models for Dynamically Tuning Data-Dependent Code
11:05	Geoff Womeldorff, LANL: Kokkos and Legion Implementations of the SNAP Proxy Application
11:15	Ryan Bleile, LLNL: Investigation of Portable Event-Based Monte Carlo Transport
11:30	Matt Martineau, UK: Investigating the Performance Portability Capabilities of OpenMP 4, Kokkos, and Raja
11:50	Leopold Grinberg, IBM: Performance Portable Single Source-Code Implementation of Sparse Linear Algebra Operations on CPUs and GPUs
12:05	Slaven Peles, LLNL: Investigating Interoperability and Performance Portability of Select LLNL Numerical Libraries
12:25	LUNCH (on your own)
1:20	**Breakout Session 1** (four separate breakouts): **Managing the Memory Hierarchy,** Doug Doerfler (LBL) and Bronson Messer (ORNL), leads **Performance Portable Abstractions,** Rob Hoekstra (SNL) and Jeff Vetter (ORNL), leads
2:50	BREAK (snacks provided)
Experience with OpenMP and Recommendations on Guiding Future Standards: Hai Ah Nam (LANL), Chair	
3:05	John Pennycook, Intel: Performance Portability of Kernel-based Abstractions
3:25	John Pennycook, Intel: Generalizing a DSL for Structured Dependency (Stencil-Like) Codes to OpenMP Loops
3:45	John Levesque, Cray: How We Can Get Hybrid OpenMP/MPI to Out-Perform All-MPI
4:05	Carlo Bertolli, IBM: Performance Portability with OpenMP on Nvidia GPUs
4:25	Jeff Larkin, NVIDIA: Performance Portability through Descriptive Parallelism
4:45	BREAK
5:15	David Appelhans, IBM: Performance Portability Experience with LLVM, OpenMP 4, and Kripke
5:30	Kevin O'Brien, IBM: OpenMP Specifications for Portability

WEDNESDAY SESSIONS, continued	
5:45	Oscar Hernandez, ORNL: Experiences with High-Level Programming Directives for Porting SPEC ACCEL on Multiple Architectures
6:00	Tom Scogland, LLNL: Performance Portability with OpenMP: Experiences with 4.5 and Looking toward 5.0
6:20	ADJOURN
6:30	Intel NDA Reception (until 20:00)

THURSDAY, APRIL 21, TECHNICAL SESSIONS

7:30a	Coffee, Pastries
8:15	Recap of Breakout #1 (each of four groups to present 8–10 minute summary)
Tools for Performance Portability and Analysis: Hai Ah Nam (LANL), Chair	
9:00	Jeanine Cook, SNL: The Importability of Performance Tools
9:10	Juan Gonzalez Garcia, IBM: Next-Gen Profiling-Infrastructure for Supercomputers Based on Hybrid Nodes
9:20	Ignacio Laguna, LLNL: STATuner—Tuning CUDA Kernels via Compiler Analysis and Machine Learning
9:35	Si Hammond, SNL: Profiling Interfaces for Parallel C++ Abstractions - KokkosP
9:50	Protonu Basu, LBL: Leveraging Compiler-Based Tools for Performance-Portability
10:10	Heidi Poxon, Cray: Adding Parallelism to HPC Applications Using Reveal
10:25	BREAK

THURSDAY SESSIONS, continued	
The Input/Output Bottleneck and Use of Burst Buffers: Mike Glass (SNL), Chair	
10:55	Mark Miller, LLNL: Probing Portable Performance of Parallel I/O Paradigms Using MACSio
11:10	Andrey Ovsyannikov, LBL: ChomboCrunch and VisIt for Carbon Sequestration and In-Transit Data Analysis Using Burst Buffers
11:30	Kathryn Mohror, LLNL: Performance Portability for Burst Buffers with the Scalable Checkpoint/Restart Library (SCR)
Use of Domain-Specific Languages for Performance Portability: Mike Glass (SNL), Chair	
11:50	David Richards, LLNL: Portable Performance in Real Applications Using Generated Code
12:05	Brian Van Straalen, LBL: AMRStencil—An Embedded DSL for Expressing Structured Adaptive Mesh Refinement Algorithms
12:20	WORKING LUNCH (provided)
1:15	**Breakout Session 2** (four separate breakouts): **OpenMP Futures,** Sriram Swaminarayan (LANL) and David Richards (LLNL), leads **Tools/Compiler/System Requirements,** Edgar Leon (LLNL) and Brian Friesen (LBL), leads
2:45	BREAK
Wrap-Up Discussions	
3:15	Recap of Breakout #2 (each of four groups to present 8–10 minute summary)
4:00	Vendor Q&A/Panel (vendor reps to discuss challenges and answer questions)
4:40	Wrap-Up/Next-Steps/Takeaways (capture follow-up goals, decide on subsequent meetings and potential topics)
5:00	ADJOURN
6:00	DINNER (provided)